

06-08-00

PTO/SB/05 MODIFIED BY AT&T CORP.

UTILITY

Attorney Docket No.

1999-0700

Total Pages

37

PATENT APPLICATION
TRANSMITTAL

First Named Inventor or Application Identifier

Samson Boodaghians

Express Mail Label No.

EL487593124US

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

ADDRESS TO:

Assistant Commissioner for Patents
Box Patent Application
Washington, D.C. 20231

1. ☒ Fee Transmittal Form
(submit an original, and a duplicate for fee processing)
2. ☒ Specification [Total Pages 21]
(preferred arrangement set forth below)

Descriptive title of invention
Cross References to Related Applications
Statement Regarding Fed sponsored R&D
Reference to Microfiche Appendix
Background of the Invention
Brief Summary of the Invention
Brief Description of the Drawings(if filed)
Detailed Description
Claim(s)
Abstract of the Disclosure

3. ☒ Drawing(s)(35 USC 113) [Total Sheets 9]
4. Oath or Declaration [Total Pages 2]

- a. ☒ Newly executed (original or copy)
b. ☐ Copy from a prior application(37 CFR 1.63(d))
(for continuation/divisional with Box 17 completed)
[Note Box 5 below]

- i. ☐ DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s)
named in the prior application, see 37 CFR
163(d)(2) and 1.33(b)

5. ☐ Incorporation by reference(usable if Box 4b is checked)
The entire disclosure of the prior application, from which a copy of the oath
or declaration is supplied under Box 4b, is considered as being part of the
disclosure of the accompanying application and is hereby incorporated by
reference herein.

6. ☐ Microfiche Computer Program (Appendix)
7. ☐ Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)
a. ☐ Computer Readable Copy
b. ☐ Paper Copy(identical to computer copy)
c. ☐ Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. ☒ Assignment Papers(cover sheet & document(s))
9. ☐ 37 CFR 3.73(b)Statement ☐ Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ Information Disclosure ☐ Copies of IDS
Statement(IDS)/PTO-1449 Citations
12. ☐ Preliminary Amendment
13. ☒ Return Receipt Postcard (MPEP 503)
(Should be specifically itemized)
14. ☐ Small Entity ☐ Statement filed in prior application,
Statement(s) Status still proper and desired
15. ☐ Certified Copy of Priority Document(s)
(if foreign priority is claimed)
16. ☐ Other :

17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior Application No:

18. CORRESPONDENCE ADDRESS

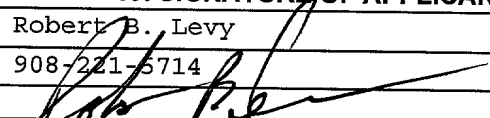
☐ Customer Number or Bar Code Label

(Insert Customer No. or Attach bar code label here)

or ☒ Correspondence address below

NAME	Samuel H. Dworetzky				
ADDRESS	AT&T CORP. P.O. Box 4110				
CITY	Middletown	STATE	New Jersey	ZIP CODE	07748-4110
COUNTRY	United States of America			FAX	732-368-6932

19. SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT REQUIRED

NAME	Robert B. Levy	Reg. #	28234	
TELEPHONE	908-221-5714			
SIGNATURE			DATE	June 7, 2000

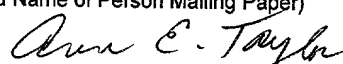
"Express Mail" Mailing Label Number EL487593124US

Date of Deposit 6/7/2000

I hereby certify that this application is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington D.C., 20231

Ann E. Taylor

(Printed Name of Person Mailing Paper)


(Signature of Person Mailing Paper)

TECHNIQUES FOR INTRODUCING IN-BAND NETWORK MANAGEMENT PACKETS IN MULTI-PROTOCOL LABEL SWITCHING NETWORKS

U.S. Patent Application Serial No. TBD (Attorney Docket No. 3493.85736),
5 entitled "Loopback Capability for Bi-Directional Multi-Protocol Label Switching Traffic
Engineered Trunks" and originally filed the same day as the present application, is hereby
incorporated by reference.

FIELD OF THE INVENTION

The present invention relates generally to network operation, administration, and
10 maintenance (OA&M) functions for communication networks. More particularly, the
present invention relates to the use of in-band network management packets and
corresponding techniques for distinguishing in-band network management packets from
user packets.

15 BACKGROUND OF THE INVENTION

A typical digital communications network has a network architecture that is based
upon the Open Systems Interconnection (OSI) Reference Model for providing
communication between a multiplicity of interconnected digital end systems or "nodes."
The OSI Reference Model divides networking protocols into seven layers, which, in
20 ascending order of abstraction, are: 1) the physical layer, 2) the data-link layer, 3) the
network layer, 4) the transport layer, 5) the session layer, 6) the presentation layer, and 7)
the application layer.

Local area networks (LANs), i.e., short-distance communications networks,
operate at layer 2 in the OSI model. Routers operate at layer 3 in the OSI model and may
25 connect two LANs or other types of networks having different protocols. More
specifically, routers at the network layer terminate local data-link layer protocols and
utilize network layer addresses and data frame restructuring to communicate with each
other.

Internet Protocol (IP) is a typical layer 3 routing protocol. For IP routing, a router
30 receives a packet and determines a next hop, i.e., the next destination for the received
packet in the path towards the final packet destination. Typically, each router in the path

to the final destination of the packet analyzes the packet header for identifying the packet's destination address and runs a routing algorithm for determining the next hop towards the identified destination address.

Multi-Protocol Label Switching (MPLS) optimizes conventional routing techniques by assigning labels to a Forwarding Equivalent Class (FEC). A FEC is defined as a set of packets that can be handled equivalently for the purpose of forwarding and thus is suitable for binding to a label. Once a binding between a FEC and a label is done, it is not necessary for each label switching router (LSR) in a label-switched path (LSP), i.e., a path through one or more LSRs followed by packets having the same FEC, to analyze a received packet's IP header for determining the packet's destination address. Instead, LSRs make forwarding decisions based on the label attached to the packet, and consequently, packets are routed through the network faster.

MPLS is being developed for high-speed networks that, for example, are used by some Internet-service providers (ISPs). MPLS is currently being standardized by the MPLS Working Group of the Internet Engineering Task Force (IETF).

FIG. 1 illustrates an MPLS packet having data-link header 8, MPLS shim header 7, IP header 6 and payload 5. MPLS uses shim header 7 located between the data-link layer header (i.e., data-link header 8) and the network layer header (i.e., IP header 6) in the MPLS packet for integrating IP routing with label switching. MPLS uses shim header 7 encapsulation shown in FIG. 1 for transporting IP Packets. MPLS can operate on any data-link layer media (e.g., ATM, FR, PPP), but MPLS currently serves only the IP client network layer. Shim header 7 generally includes a series of label stack entries. As shown in FIG. 1, a label stack entry contains a 20-bit label field 1, a 3-bit experimental field 2, a single bit field 3 indicating the bottom of the label stack and an 8-bit time-to-live (TTL) field 4.

Similar to conventional routing table entries, each LSR in a MPLS network may include a forwarding table having next hop label forwarding entries (NHLFEs). Each NHLFE, among other information, contains the physical interfaces or ports, the incoming label, and the outgoing label for the next hop for a received packet. A label in a label stack entry of a received packet is used as an index for retrieving a NHLFE containing the next hop for the received packet. Generally, the label from the label stack entry on the

top of the label stack is used to index the NHLFEs in the LSR. After identifying the NHLFE for the next hop, the outgoing label for the next hop, which is retrieved from the identified NHLFE, is placed on top of the label stack for the packet, and the packet is transmitted to the next hop. This label switching technique is used by the LSRs for routing MPLS packets through the MPLS network.

FIG. 2 illustrates a reference topology for a typical MPLS network 45. MPLS network 45 includes a set of nodes or LSRs for performing MPLS routing and forwarding. The LSRs in MPLS network 45 include intermediate LSRs and label-edge routers, e.g., LER 1 and LER 2. The set of contiguous nodes that an MPLS packet traverses in the MPLS network is called a Label Switched Path (LSP).

MPLS network 45 includes a bi-directional traffic engineering trunk (BTT) 42 having two traffic engineering trunks with the same endpoints, i.e., LER 1 and LER 2, and opposite directions of transmission. The two traffic trunks that form BTT 42 traverse two unidirectional explicitly routed label-switched paths (ER-LSPs) between LER 1 and LER 2. . An ER-LSP is a LSP defined by a management system or a single LSR that is typically the ingress LSR for that particular direction of transmission, e.g., LER 1 or LER 2. ER-LSPs are set up independent of IP shortest path routing algorithms. BTT 42 is formed of two traffic trunks traversing two ER-LSPs in opposite directions. One traffic trunk flows downstream from ingress LER 1 towards egress LER 2 on one ER-LSP. The other traffic trunk flows upstream from egress LER 2 towards ingress LER 1 on the other ER-LSP. Consequently, BTT 42 traverses a complete round-trip path between LER 1 and LER 2.

It is envisioned that ER-LSPs, which form BTTs, will be set up for carrying possibly hundreds to several thousands of individual IP flows. Hence, it is crucial to ascertain parameters of a BTT, such as connectivity, delay and other quality of service (QoS) parameters that may effect traffic flow in the BTT.

An in-band network management packet (INMP) is a packet that carries operation, administration and maintenance (OA&M) information. INMPs can be used for testing parameters of a BTT and sending OA&M commands to an LSR. U.S. Patent Application Serial No. TBD (Attorney Docket No. 3493.85736) describes two types of

5

15

25

hop without the need to read payload 5. However, if the packet is an INMP, the receiving LSR must read and process OA&M information in the INMP payload.

Hence, a need exists for providing LSRs with the capability to determine whether an MPLS packet is an INMP or a user packet.

5

SUMMARY OF THE INVENTION

It is an aspect of the present invention to use INMPs for conveying OA&M information in an MPLS network. It is another aspect of the present invention to use different fields of the MPLS header for distinguishing INMPs from user data.

10 In accordance with the present invention, there is provided a system and method for distinguishing an INMP packet from a user packet. In accordance with the present invention, a predetermined 3-bit code in the EXP field of the MPLS header is used to determine whether the received packet is an INMP or a user packet. In another preferred embodiment of the present invention, a predetermined 1-bit code in the EXP field of the
15 MPLS header is used to determine whether the received packet is an INMP or a user packet.

Alternatively, in accordance with the present invention, a predetermined code in the TTL field is used to determine whether the received packet is an INMP or a user packet.

20 Alternatively, in accordance with the present invention, a label (i.e., a reserved label) is designated and used to determine whether a received packet is an INMP or a user packet. Using a reserved label requires modification to conventional label switching techniques, because the reserved label is inserted in the top layer of the shim header label stack where the switching label (i.e., the label used to determine the next hop) is normally
25 located. Thus, in accordance with, the present invention, a reserved label is inserted on top of a label stack in a shim header in a packet, and a label below the reserved label on the label stack is used to determine the next hop for the packet. Alternatively, the reserved label is inserted in a shim header of a packet, whereas the switching label is inserted in a designated field in the INMP payload.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the accompanying figures in which like reference numerals indicate similar elements and in which:

- FIG. 1 is a diagram illustrating a format of a packet header used in a typical MPLS network;
- FIG. 2 is a schematic block diagram illustrating the topology of a typical MPLS network;
- FIG. 3 is a diagram of an EXP field in an MPLS shim header;
- FIG. 4 is a schematic block diagram of a label switched router of the present invention;
- FIG. 5 is a diagram of a TTL field in an MPLS shim header;
- FIG. 6 is a flow diagram of a preferred embodiment of the present invention that uses an INMP ID bit;
- FIG. 7 is a flow diagram of a preferred embodiment of the present invention that uses an INMP ID code;
- FIG. 8 is a flow diagram of a preferred embodiment of the present invention that uses the TTL field;
- FIG. 9 is a flow diagram of a preferred embodiment of the present invention that uses a reserved label; and
- FIG. 10 is a flow diagram of another preferred embodiment of the present invention that uses a reserved label.

DETAILED DESCRIPTION OF THE INVENTION

1. Experimental Field (or EXP bits)

The EXP field 2 in shim header 7 of FIG. 1 has a length of three bits. FIG. 3 illustrates EXP field 2 having three bits A-C, bit A being the most significant bit and bit

C being the least significant bit. Currently, the use for the EXP bits is not standardized, but there is a proposal to use the EXP bits for signaling class of service (CoS). For example, an LSR includes a plurality of queues for temporarily storing received MPLS packets that need to be routed. The MPLS packets are placed on different queues as they
 5 are ready to exit a port on the LSR. MPLS packets in the higher priority queue (e.g., the queue designated for high CoS packets) are served before the remaining packets in other queues, which contain packets of lower CoS. Alternatively, each queue can be given a share of the LSR port and line bandwidth, which is proportional to its CoS designation as compared to other queues. Since there are three bits in EXP field 2, up to eight classes
 10 can be designated.

In a preferred embodiment of the present invention, a bit in EXP field 2, hereinafter referred to as an INMP ID bit, is designated to carry coded information for differentiating INMPs from user packets. For example, bit A in EXP field 2 is designated as the INMP ID bit, and an INMP ID bit having a value of "1" identifies an INMP. One
 15 of ordinary skill in the art, however, can readily designate any one of bits A-C as the INMP ID bit, and an INMP ID bit having a value of "0" may also be used for identifying an INMP. An LSR receiving an MPLS packet with bit A in EXP field 2 having a value of "1" determines from bit A that the received MPLS packet is an INMP. After making this determination, the LSR processes the INMP to, e.g., test a parameter of the BTT or
 20 perform a command, as discussed above.

In another preferred embodiment, a predetermined 3-bit code carried in EXP field 2, hereinafter referred to as an INMP ID code, identifies an INMP. For example, "111" in EXP field 2 is designated as the INMP ID code, and an LSR receiving an MPLS packet having "111" in EXP field 2 determines that the received packet is an INMP. One
 25 of ordinary skill in the art, however, can readily designate any 3-bit code for the INMP ID code.

As discussed above, there is a proposal to use the EXP bits for designating CoS for the MPLS packet. Designating an INMP ID code limits the number of classes that may be designated from eight to seven, because one of the eight codes is allocated as an
 30 INMP ID code. Also, designating an INMP ID bit limits the number of classes that may be designated from eight to four, because one bit from EXP field 2 is allocated as an

INMP ID bit. In many networks, however, up to four classes may be adequate, so allocating a code or a bit in EXP field 2 for identifying an INMP can be acceptable.

FIG. 4 is a schematic block diagram of a preferred embodiment of an LSR in MPLS network 40. The LSR shown in FIG. 4 includes ports 50-53, processing circuitry 65 and switching fabric 60. Each of ports 50-53 includes transmitting circuitry, receiving circuitry and packet assembly circuitry, as is known in the art. Ports 50-53 are connected to processing circuitry 65 through switching fabric 60. The switching fabric, for example, may be a high-speed bus and include one or more multiplexors and de-multiplexors. Processing circuitry 65 includes a processor 61 and memory 62. Processor 61 may include a high-speed processor and performs forwarding/routing functions. Memory 62 may include, among other information, NHLFEs stored, for example, in a table or database for determining a next hop.

Typical operation of the LSR shown in FIG. 4 may include receiving incoming packets at ports 50-53. Packet information from the incoming packets is sent to processing circuitry 65 through switching fabric 60. Processing circuitry 65 processes the packet information to determine the next hop for an incoming packet. For example, processing circuitry 65 can identify an incoming label for each packet, and retrieve an NHLFE corresponding to the incoming label from a table stored in memory 62. The retrieved NHLFE contains the label for the next hop. Then processing circuitry 65 forwards the packet information with the new label through switching fabric 60 to the port associated with the next hop. The packet information is assembled into an outgoing packet. The outgoing packet is then transmitted to the next hop.

FIG. 6 illustrates a flow diagram of the preferred embodiment of the present invention that uses an INMP ID bit for identifying an INMP. In step 100, LER 1 constructs an MPLS INMP having a payload for conveying OA&M information, and LER 1 inserts the INMP ID code in shim header 7. Any LER, such as LER 2, however, may construct an MPLS INMP. In step 110, LER 1 transmits an MPLS packet on a traffic trunk, such as BTT 42, to LSR 1, i.e., the next hop on BTT 42 downstream. In step 120, LSR 1 receives the MPLS packet on, for example, port 50. The packet information is forwarded to processing circuitry 65. In step 130, processing circuitry 65 reads the MPLS packet header, including EXP field 2 in shim header 7. In step 140,

processing circuitry 65 identifies the INMP ID bit from EXP field 2 and compares the INMP ID bit with a predetermined code, such as "0" or "1". If the INMP ID bit equals the predetermined code, the processing circuitry identifies the packet as an INMP. If the INMP ID bit equals the predetermined code, processing circuitry 65 processes the INMP.

- 5 In step 100, the MPLS packet is constructed as an INMP, and thus, in step 150, that INMP is processed. Processing the INMP, for example, includes determining whether the INMP is a test INMP or a command INMP. If the INMP is a test INMP, processing the INMP, for example, includes testing a parameter of the BTT and transmitting the INMP to the next hop. If the INMP is a command INMP, processing the INMP, for
- 10 example, includes performing the command and terminating the INMP if the command is designated for that LSR. If the INMP ID bit is not equal to the predetermined code, in step 160, processing circuitry 65 determines a next hop, using a NHLFE, for the received packet and the MPLS packet is label-switched to LSR 2, i.e., the next hop on BTT 42.

- FIG. 7 illustrates a flow diagram of the preferred embodiment of the present invention that uses an INMP ID code for identifying an INMP. The steps are
- 15 substantially the same as those shown in FIG. 6 for the preferred embodiment using an INMP ID bit, except the INMP ID code is compared to a predetermined code during step 140.

20 2. Time-to-Live Field

- Currently, as shown in FIG. 1, the one octet TTL field 4 in MPLS shim header 7 is used in the same manner as the TTL field in IP header 6. That is, the IP TTL field is copied into the MPLS header TTL field upon a packet's entry into MPLS network 45 of
- FIG. 2. The MPLS header TTL field is then decremented by one as the MPLS packet
- 25 traverses each LSR towards its final destination. If the TTL field becomes zero, the packet is discarded since that is a strong indication that the MPLS packet is routed in an endless routing loop among LSRs in MPLS network 45.

- In another preferred embodiment of the present invention, TTL field 4 of FIG. 1 can be used to code information to distinguish INMPs from user packets. In this
- 30 preferred embodiment, a single bit from TTL field 4, hereinafter referred to as a TTL ID bit, is designated as a flag to differentiate an INMP from a user packet, similarly to an

INMP ID bit in EXP field 2. TTL field 4 is shown in FIG. 5 with A1 as the most significant bit. Bit A1 in TTL field 4 is designated as the TTL ID bit, and the TTL ID bit having a value of "1" identifies an INMP. One of ordinary skill in the art, however, can readily designate any one of bits A1-H1 as the TTL ID bit, and a TTL ID bit having a value of "0" may also be used for identifying an INMP. An LSR receiving an MPLS packet with bit A1 in TTL field 4 having a value of "1" determines from bit A1 that the received packet is an INMP. After determining that the packet is an INMP, the LSR processes the INMP.

Designating a bit in the TTL field for distinguishing INMPs from user data limits the number of hops a packet can traverse to 128 hops (i.e., 2 to the power of 7) before it is discarded. However, 128 hops is a more than adequate threshold for discarding packets routed in an endless loop. Also, similar to the INMP ID code, an 8-bit TTL code may be designated and used to determine whether a received packet is an INMP or a user packet.

For Hop-by-Hop MPLS, where the LSPs are formed using IP routing algorithms, the TTL field is necessary, because the potential for routing loops exists. For explicitly routed LSPs (ER-LSPs) the routes are calculated independent of IP routing algorithms, and thus the TTL field may not be necessary since the potential for routing loops can be eliminated. However, the network will still have to keep track of the TTL field as if the packet has traversed an IP network. In other words, when the packet emerges from the MPLS domain and into another MPLS or IP domain, it must contain the correct TTL value. The situation where the TTL field is not available to decrement occurs in ATM and FR MPLS networks, because ATM and FR layer 2 headers do not include a TTL field. In such cases, the Label Distribution Protocol (LDP) utilizes the Hop-Count variable to pre-calculate the hop count associated with an LSP. At the egress LER, this Hop-Count is then subtracted from the original TTL value of the packet (i.e., the TTL value of the packet upon entry into the MPLS domain). When the INMP ID code method is used to identify an INMP, the TTL field is not available to decrement. Therefore, the aforementioned Hop-Count method can be used for processing the TTL field upon exiting the MPLS domain.

FIG. 8 illustrates a flow diagram of the preferred embodiment of the present invention that uses a TTL ID bit for identifying an INMP. In step 300, LER 1 constructs an MPLS INMP having a payload for conveying OA&M information and inserts the TTL ID bit in shim header 7. Any LER, however, may construct the MPLS INMP, such as LER 2. In step 310, LER 1 transmits an MPLS packet on a traffic trunk, such as BTT 42 of FIG. 4 to LSR 1, i.e., the next hop on BTT 42 downstream. In step 320, LSR 1 receives the MPLS packet on, for example, port 50 of FIG. 4. The packet information is forwarded to processing circuitry 65 of FIG. 4. In step 330, processing circuitry 65 reads the MPLS packet header, including TTL field 4 in shim header 7. In step 340, processing circuitry 65 identifies the TTL ID bit from TTL field 4 and compares the TTL ID bit with a predetermined code, such as "0" or "1". If the TTL ID bit equals the predetermined code, the processing circuitry identifies the packet as an INMP. If the TTL ID bit equals the predetermined code, in step 350, processing circuitry 65 processes the INMP. In step 300, the constructed MPLS packet is an INMP, and thus, in step 350, the INMP is processed. Processing the INMP, for example, includes determining whether the INMP is a test INMP or a command INMP. If the INMP is a test INMP, processing the INMP, for example, includes testing a parameter of the BTT and transmitting the INMP to the next hop. If the INMP is a command INMP, processing the INMP, for example, includes performing the command and terminating the INMP if the command is designated for that LSR. If the INMP ID bit is not equal to the predetermined code, in step 360, processing circuitry 65 determines that the packet is a user packet. Then, processing circuitry 65 determines the next hop, using a NHLFE, and the MPLS packet is label-switched to its next hop.

25 3. Reserved Label

An MPLS packet is transported on an LSP using the existing label-switching technique described above. For example, an LSR in network 45 receiving an MPLS packet determines the next hop for the received packet using a label from the top of the label stack. A label, however, may be reserved for identifying an INMP (i.e., distinguishing an INMP from a user packet). The reserved label is a specific label that

identifies an INMP, and thus the reserved label can be used to distinguish an INMP from a user packet.

Using the reserved label requires some enhancements to the existing label-switching methods to support transport of INMPs. In a preferred embodiment of the present invention, a label stack entry including a reserved label is inserted on top of the label stack, and the switching label is inserted in a label stack entry directly below the reserved label. Once an LSR identifies the reserved label, the LSR maintains the reserved label on top of the label stack, but the LSR switches labels using the switching label in the label stack entry immediately below the reserved label. Therefore, the LSR determines the next hop for the packet using a label directly below the label (i.e., the reserved label) on top of the label stack instead of using the label on top of the label stack.

FIG. 9 illustrates the preferred embodiment of the present invention that maintains the reserved label on top of the label stack. In step 400, ingress LER 1 constructs an INMP having an INMP payload for conveying OA&M information. LER1 then inserts the reserved label on top of the label stack in shim header 7. Note that the label normally used for switching (i.e., the switching label) now resides immediately below the aforementioned reserved label. In step 410, LER 1 transmits the INMP on BTT 42 to LSR 1, i.e., the next hop downstream on BTT 42, using the switching label. Any LER, however, may construct the INMP, and the INMP may be transmitted on a unidirectional traffic trunk as well as a bi-directional traffic trunk. In step 420, LSR 1 receives the INMP on, for example, port 50. The packet information is forwarded to processing circuitry 65. In step 430, processing circuitry 65 reads the label from the top of the label stack in shim header 7. In step 440, processing circuitry 65 determines whether the label is the reserved label. If the label read from the top of the label stack is the reserved label, the processing circuitry identifies the received packet as an INMP. In step 450, if the label read from the top of the label stack is the reserved label, processing circuitry 65 processes the INMP. In step 400, the constructed MPLS packet is an INMP, and thus, in step 450, the INMP is processed. Processing the INMP, for example, may include determining whether the INMP is a test INMP or a command INMP and transmitting the INMP to a next hop if the INMP is a test INMP or performing the command and

terminating the INMP if the INMP is a command INMP. In step 450, if the INMP is not terminated, processing circuitry 65 determines the next hop for the received packet using a label in a label stack entry directly below the reserved label, which is on top of the label stack. The INMP is then transmitted to the identified next hop using the label directly
 5 below the reserved label. In step 460, if the label on top of the label stack is not the reserved label, processing circuitry 65 determines the next hop for the received packet using the label from the top of the label stack.

In another preferred embodiment of the present invention, a reserved label is carried in a label stack entry in shim header 7, and a field called the “switching label
 10 field” in the INMP payload is designated for carrying the switching label. In this embodiment, like the previously described embodiment, the reserved label is carried in a label stack entry in shim header 7. However, in this preferred embodiment, the LSR receiving the INMP determines the next hop using a switching label carried in a field in the INMP payload of the received INMP. For example, the INMP payload may be
 15 partitioned into a plurality of fields, and one of the fields may be designated for carrying the switching label. Currently, the protocol and semantic structure of an INMP payload is not determined, and one of ordinary skill in the art may designate a field for carrying the switching label when the protocol and complete semantic structure of the INMP payload is standardized.

FIG. 10 illustrates the preferred embodiment of the present invention that carries the switching label in a field in an INMP payload without regard to its position in the stack in contrast to FIG. 9. In step 500, ingress LER 1 constructs an INMP having the reserved label in a label stack entry in shim header 7. In step 510, LER 1 transmits the INMP on BTT 42 to LSR 1, i.e., the next hop on BTT 42 downstream. Any LER,
 25 however, may construct the INMP and the INMP may be transmitted on a unidirectional as well as a bi-directional traffic trunk. In step 520, LSR 1 receives the INMP on, for example, port 50 of FIG. 4. The packet information is forwarded to processing circuitry 65 of FIG. 4. In step 530, processing circuitry 65 reads the label from shim header 7. In step 540, processing circuitry 65 determines whether the switching label is the reserved
 30 label. If the label read from a shim header 7 is the reserved label, processing circuitry 65 identifies the packet as an INMP and processes the INMP during step 550. In step 500,

the constructed MPLS packet is an INMP, and thus, in step 550, the INMP is processed. Processing the INMP, for example, includes determining whether the INMP is a test INMP or a command INMP. If the INMP is a test INMP, processing the INMP, for example, includes testing a parameter of the BTT and transmitting the INMP to the next hop. If the INMP is a command INMP, processing the INMP, for example, includes performing the command and terminating the INMP if the command is designated for that LSR. If the INMP is not terminated, during step 550 processing circuitry 65 determines the next hop for the received packet using a switching label in a field in the INMP payload and the reserved label is maintained in shim header 7. Processing circuitry 65 can either switch the switching label in this field with the outgoing label identified by the retrieved NHFLE, or processing circuitry 65 can append the outgoing label to the previously appended switching label(s) in the designated field in the INMP payload. With the latter approach, a label trace of the INMP is maintained in the designated field of the INMP payload. In step 560, if the label in shim header 7 is not the reserved label, processing circuitry 65 determines the next hop for the received packet using the label on top of the label stack in shim header 7.

The present invention is applicable to Internet backbones and enterprise networks which use MPLS as a transport mechanism. In addition, some aspects of the present invention can be implemented for ATM, Frame Relay, and optical networks utilizing label-switching techniques.

What has been described are the preferred embodiments of the present invention. It will be apparent, however, to those skilled in the art that it is possible to embody the invention in specific forms other than those disclosed in the preferred embodiments described above. This may be done without departing from the spirit of the invention, and the preferred embodiments are merely illustrative and should not be considered restrictive in any way. The scope of the invention is given by the appended claims, rather than the preceding description.

What is claimed

- 1 1. A method of introducing in-band network management packets in a network
2 comprising steps of:
3 constructing a packet including a header;
4 inserting a predetermined code in a field in the header; and
5 determining whether the packet includes an in-band network management packet
6 or a user packet using the predetermined code.
- 1 2. The method of claim 1, wherein the field for inserting the predetermined code is an
2 experimental field.
- 1 3. The method of claim 2, wherein the predetermined code is a three-bit code.
- 1 4. The method of claim 3, wherein the predetermined code is a one-bit code.
- 1 5. The method of claim 1, wherein the field for inserting the predetermined code
2 indicates class of service for the packet.
- 1 6. The method of claim 2, wherein the field for inserting the predetermined code is a
2 time-to-live field.
- 1 7. The method of claim 6, wherein the predetermined code is a one-bit code.
- 1 8. The method of claim 1, wherein the constructed packet is a multi-protocol label-
2 switching packet.
- 1 9. The method of claim 1, wherein the header includes a shim header, and the field
2 wherein the predetermined code is inserted is located in the shim header.
- 1 10. The method of claim 1, further including a step of:
2 transmitting the constructed packet on a multi-protocol label switching network.

1 11. A method of introducing in-band network management packets in a network,
 2 comprising a step of:
 3 determining whether a packet is an in-band network management packet or a user
 4 packet.

1 12. The method of claim 11, wherein the step of determining whether a packet is an in-
 2 band network management packet or a user packet further includes:
 3 using a predetermined code to distinguish an in-band network management packet
 4 from a user packet.

1 13. The method of claim 12, wherein the packet includes a shim header and the
 2 predetermined code is inserted in an experimental field located in the shim header.

1 14. The method of claim 12, wherein the packet includes a shim header and the
 2 predetermined code is inserted in a time-to-live field located in the shim header.

1 15. The method of claim 11, wherein the packet is multi-protocol label switching packet.

1 16. A method of introducing in-band network management packets in a network,
 2 comprising steps of:
 3 designating a label that distinguishes an in-band network management packet
 4 from a user packet;
 5 constructing a packet; and
 6 determining whether the constructed packet is an in-band network management
 7 packet or a user packet using the designated label.

1 17. The method of claim 16, wherein the constructed packet includes a header and a
 2 payload, the header including a shim header, and further including a step of:
 3 inserting the designated label in the shim header.

- 1 18. The method of claim 17, further including steps of:
2 inserting the designated label on top of a label stack in the shim header; and
3 determining a next hop for the packet using a label on the label stack below the
4 designated label.
- 1 19. The method of claim 16, wherein the packet is a multi-protocol label switching
2 packet.
- 1 20. The method of claim 17, further including steps of:
2 constructing an in-band network management packet having a payload; and
3 determining a next hop for the packet using a label in a designated field in the
4 payload of the in-band network management packet.
- 1 21. The method of claim 16, wherein the step of determining whether the constructed
2 packet is an in-band network management packet or a user packet is performed by a
3 router in a multi-protocol label switching network receiving the constructed packet..
- 1 22. A network comprising:
2 an originating router constructing an in-band network management packet; and
3 a receiving router that receives a packet and determines whether the packet is an
4 in-band network management packet or a user packet.
- 1 23. The network of claim 22, wherein the originating router inserts a predetermined code
2 in a header in the in-band network management packet, and the predetermined code
3 identifies an in-band network management packet.
- 1 24. The network of claim 23, wherein the header includes a shim header, and the
2 predetermined code is inserted in an experimental field in the shim header.
- 1 25. The network of claim 24, wherein the predetermined code is any one of a three-bit
2 code and a one-bit code.

- 1 26. The network of claim 23, wherein the header includes a shim header, and the
2 predetermined code is inserted in a time-to-live field in the shim header.
- 1 27. The network of claim 22, wherein the constructed packet is a multi-protocol label
2 switching packet.
- 1 28. The network of claim 22, wherein the network is a multi-protocol label switching
2 network.
- 1 29. The network of claim 22, wherein the originating router inserts a reserved label in a
2 header in the packet, and the receiving router uses the reserved label to determine
3 whether the packet is an in-band network management packet or a user packet.
- 1 30. A network comprising:
2 an originating router constructing an in-band network management packet and
3 inserting a reserved label in a header in the packet; and
4 a receiving router that receives a packet and determines whether the packet is an
5 in-band network management packet or a user packet using the reserved label.
- 1 31. The network of claim 30, wherein the header includes a shim header, the reserved
2 label is inserted on top of a label stack in the shim header and the receiving router
3 determines a next hop for the packet using a label on the label stack below the
4 reserved label.
- 1 32. The network of claim 30, wherein the originating router constructs an in-band
2 network management packet and the receiving router determines a next hop for the
3 packet using a label in a designated field in a payload of the constructed in-band
4 network management packet.
- 1 33. The network of claim 30, wherein the constructed packet is a multi-protocol label
2 switching packet.

- 1 34. The network of claim 30, wherein the network is a multi-protocol label switching
2 network.
- 1 35. A router comprising:
2 reception circuitry that receives an incoming packet; and
3 processing circuitry that identifies a predetermined code and determines whether
4 the incoming packet is an in-band network management packet or a user packet using
5 the predetermined code.
- 1 36. The router of claim 35, wherein the processing circuitry identifies the predetermined
2 code from an experimental field in a shim header of the received packet.
- 1 37. The router of claim 35, wherein the predetermined code is any one of a one-bit and
2 three-bit code.
- 1 38. The router of claim 35, wherein the processing circuitry identifies the predetermined
2 code from a time-to-live field in a shim header of the received packet.
- 1 39. The router of claim 35, wherein the constructed packet is a multi-protocol label
2 switching packet.
- 1 40. The router of claim 35, wherein the network is a multi-protocol label switching
2 network.
- 1 41. A router comprising:
2 reception circuitry that receives an incoming packet having a header that includes
3 a shim header and a payload; and
4 processing circuitry that identifies a reserved label in the shim header in the
5 packet and determines whether the incoming packet is an in-band network
6 management packet or a user packet using the reserved label.

- 1 42. The router of claim 41, wherein the reserved label is on top of a label stack in the
2 shim header and the processing circuitry determines the next hop for the incoming
3 packet using a label below the reserved label on the label stack.
- 1 43. The router of claim 41, wherein the processing circuitry determines a next hop for the
2 incoming packet using a label in a designated field in a payload of an in-band
3 network management packet.
- 1 44. The router of claim 41, wherein the incoming packet is a multi-protocol label
2 switching packet.
- 1 45. The router of claim 41, wherein the router is a multi-protocol label switching router.

ABSTRACT

A system and method is provided for introducing in-band network management packets (INMPs) in a Multi-Protocol Label Switching (MPLS) network. This system and method of the invention uses INMPs for carrying OA&M information to label switching routers (LSRs) for effectively managing and operating MPLS-based networks. This system and method of the invention also includes techniques for distinguishing INMPs from user packets in an MPLS network. The system and method of the invention further includes using a predetermined code in a shim header of an MPLS packet to determine whether an MPLS packet is an INMP or a user packet. The predetermined code may be provided in an experimental field or a time-to-live field in the shim header of the packet. Alternatively, a label may be reserved for distinguishing an INMP from a user packet.




FIG. 1

PRIOR ART

007090" 99168560




FIG. 2

Prior Art




FIG. 3




FIG. 5




FIG. 4

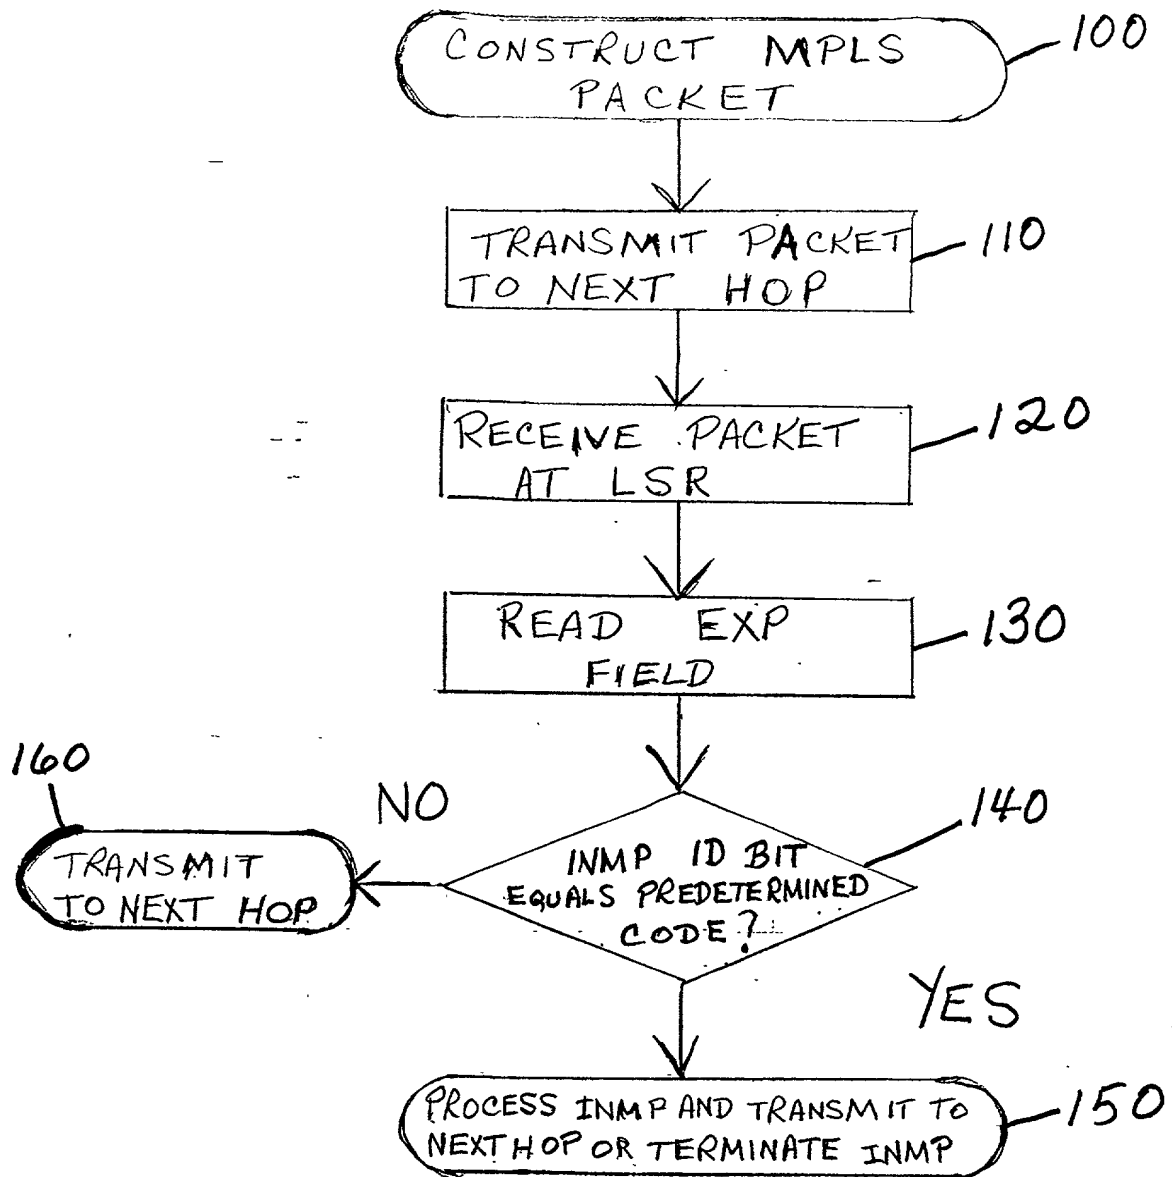


FIG. 6

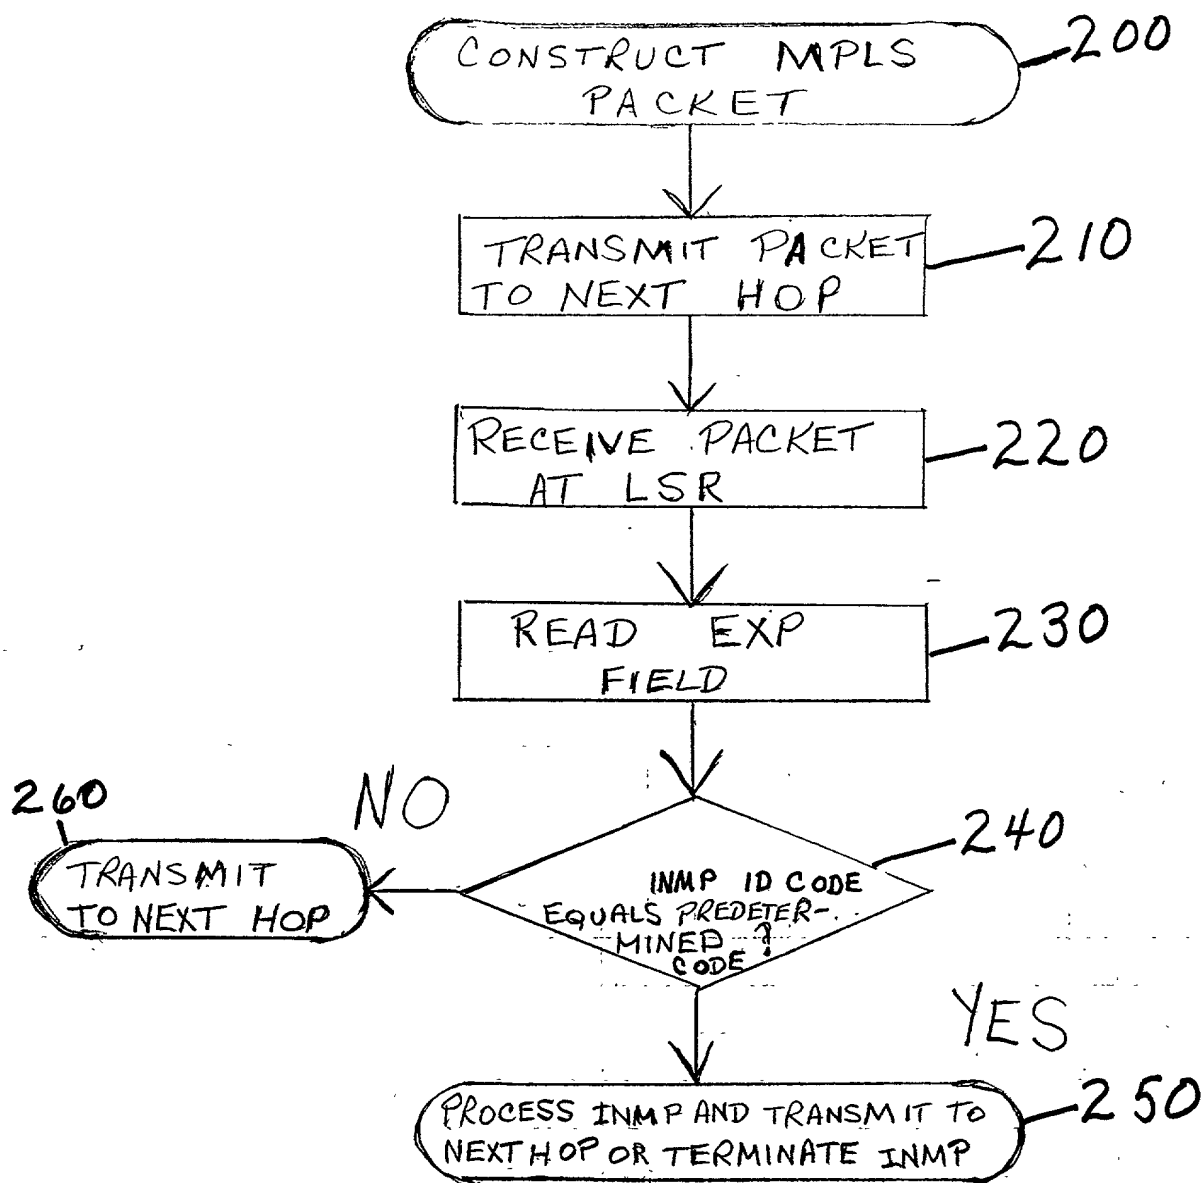


FIG. 7

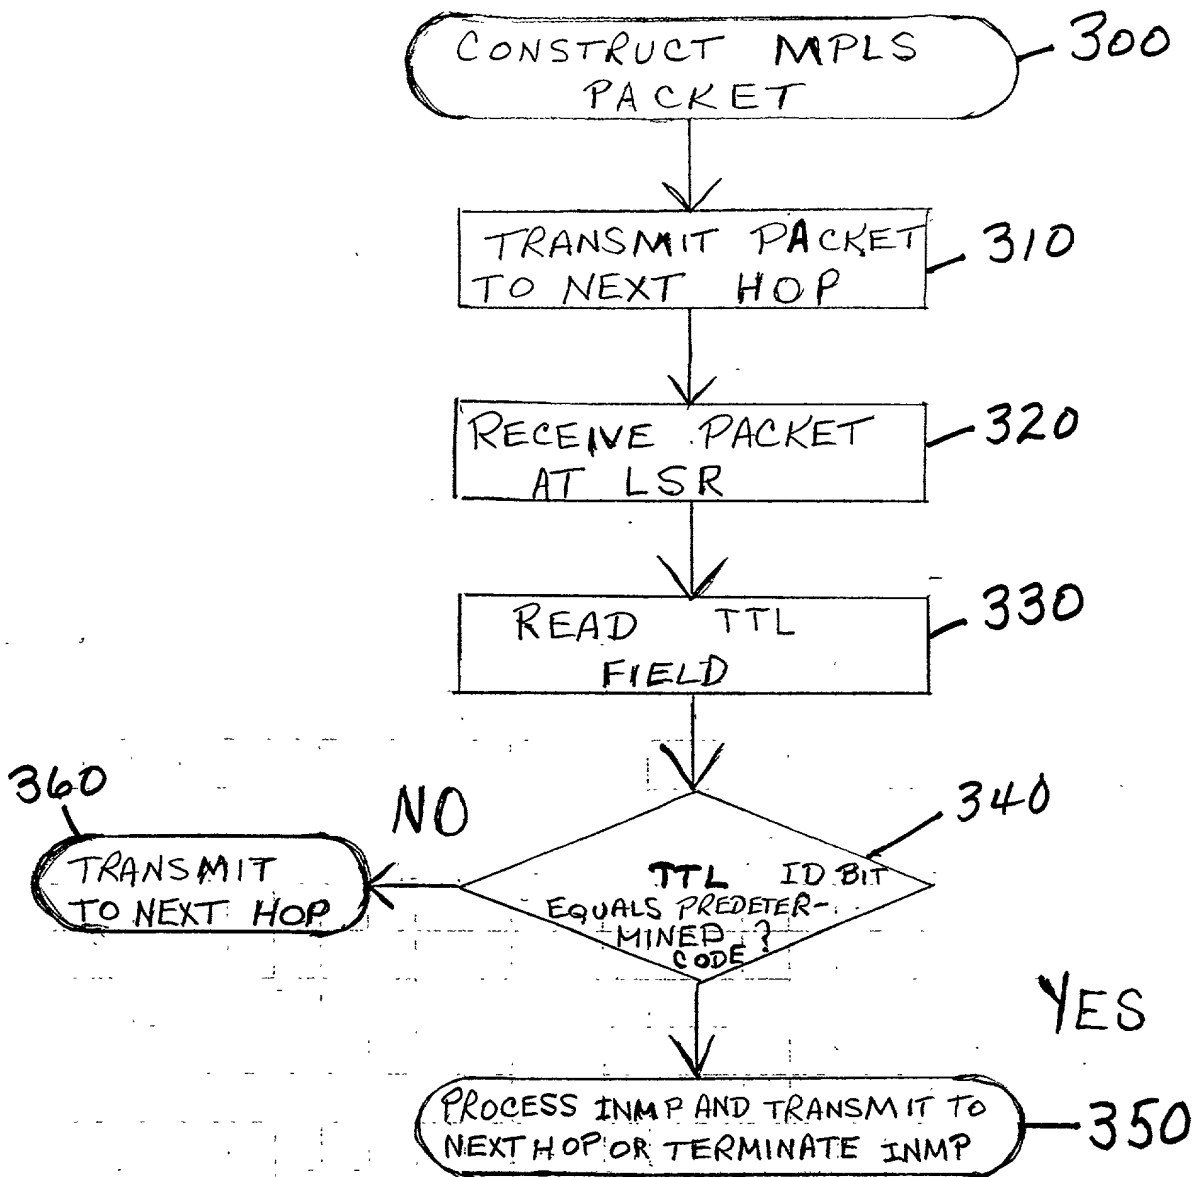
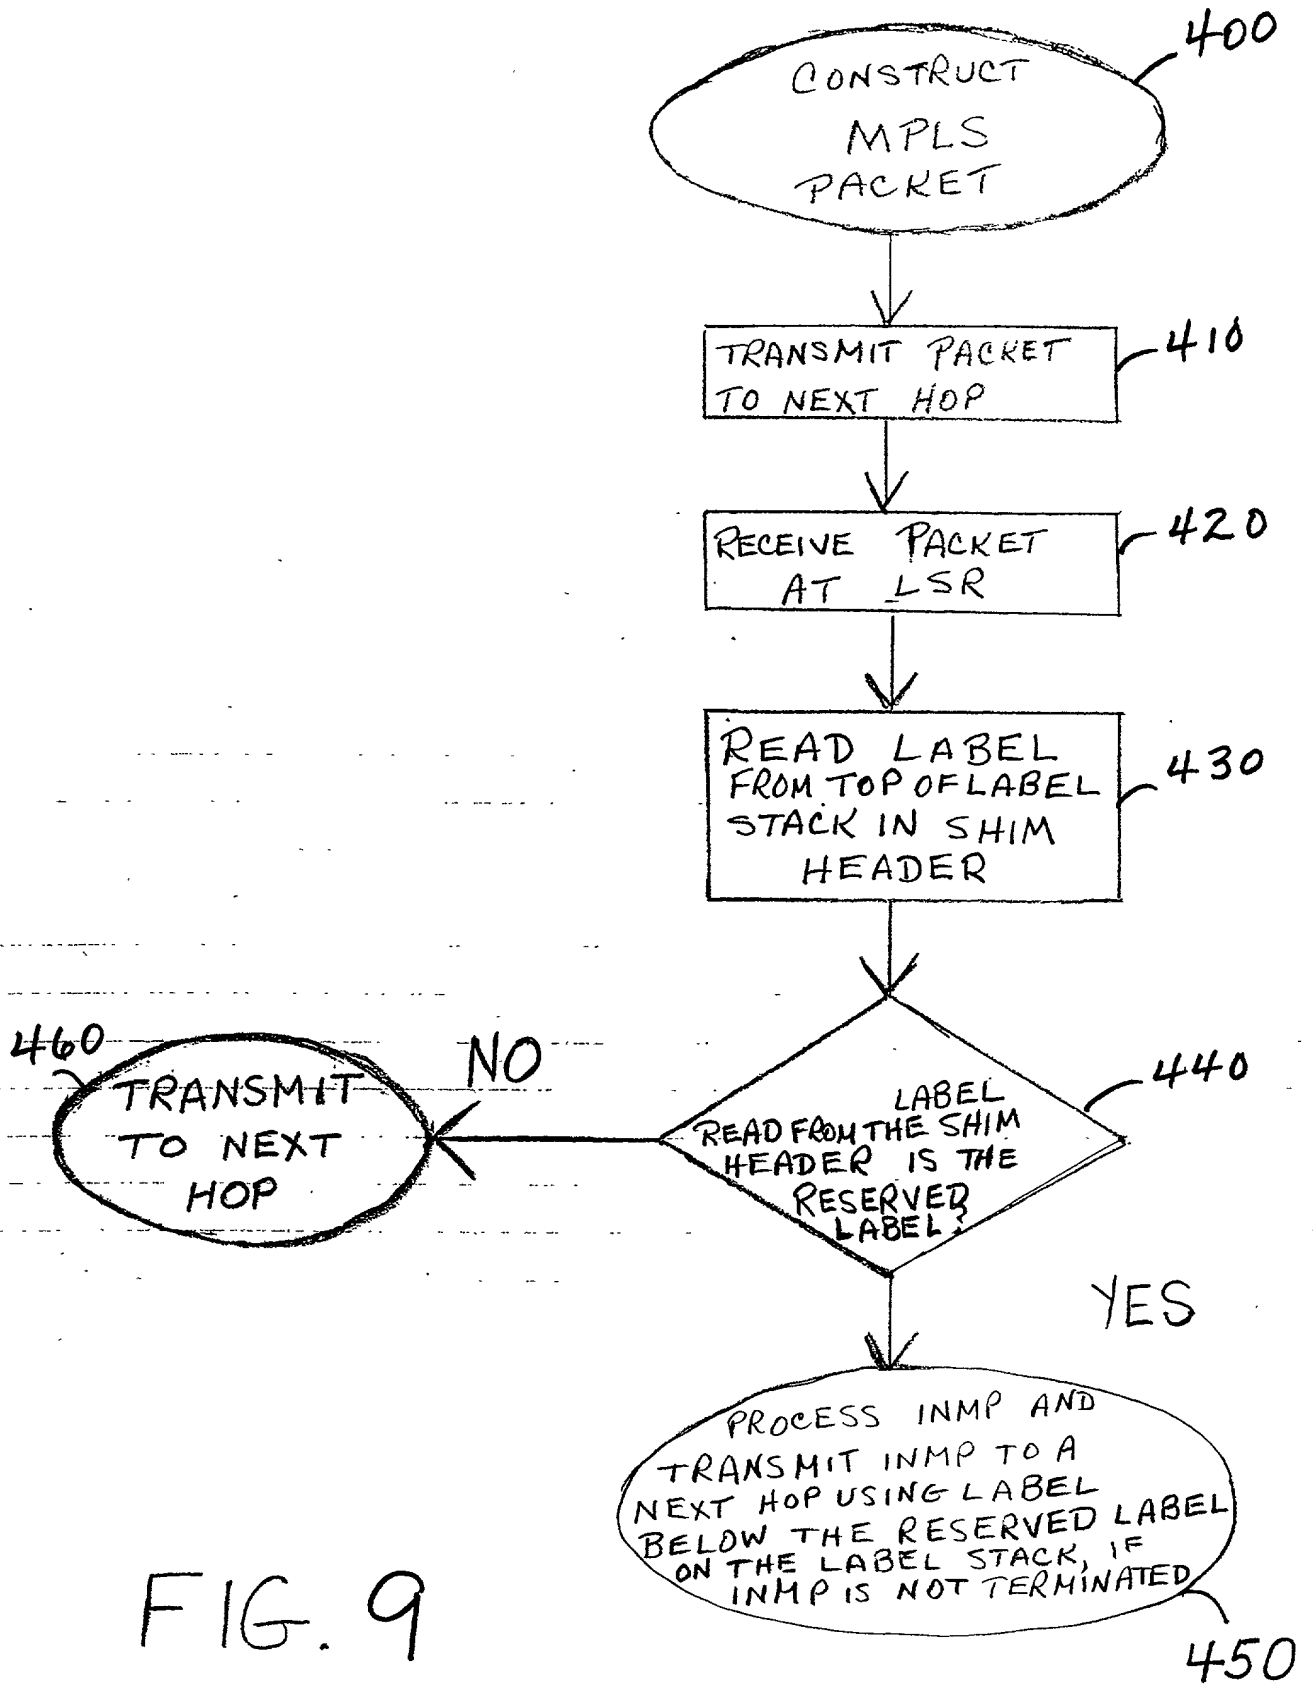


FIG. 8



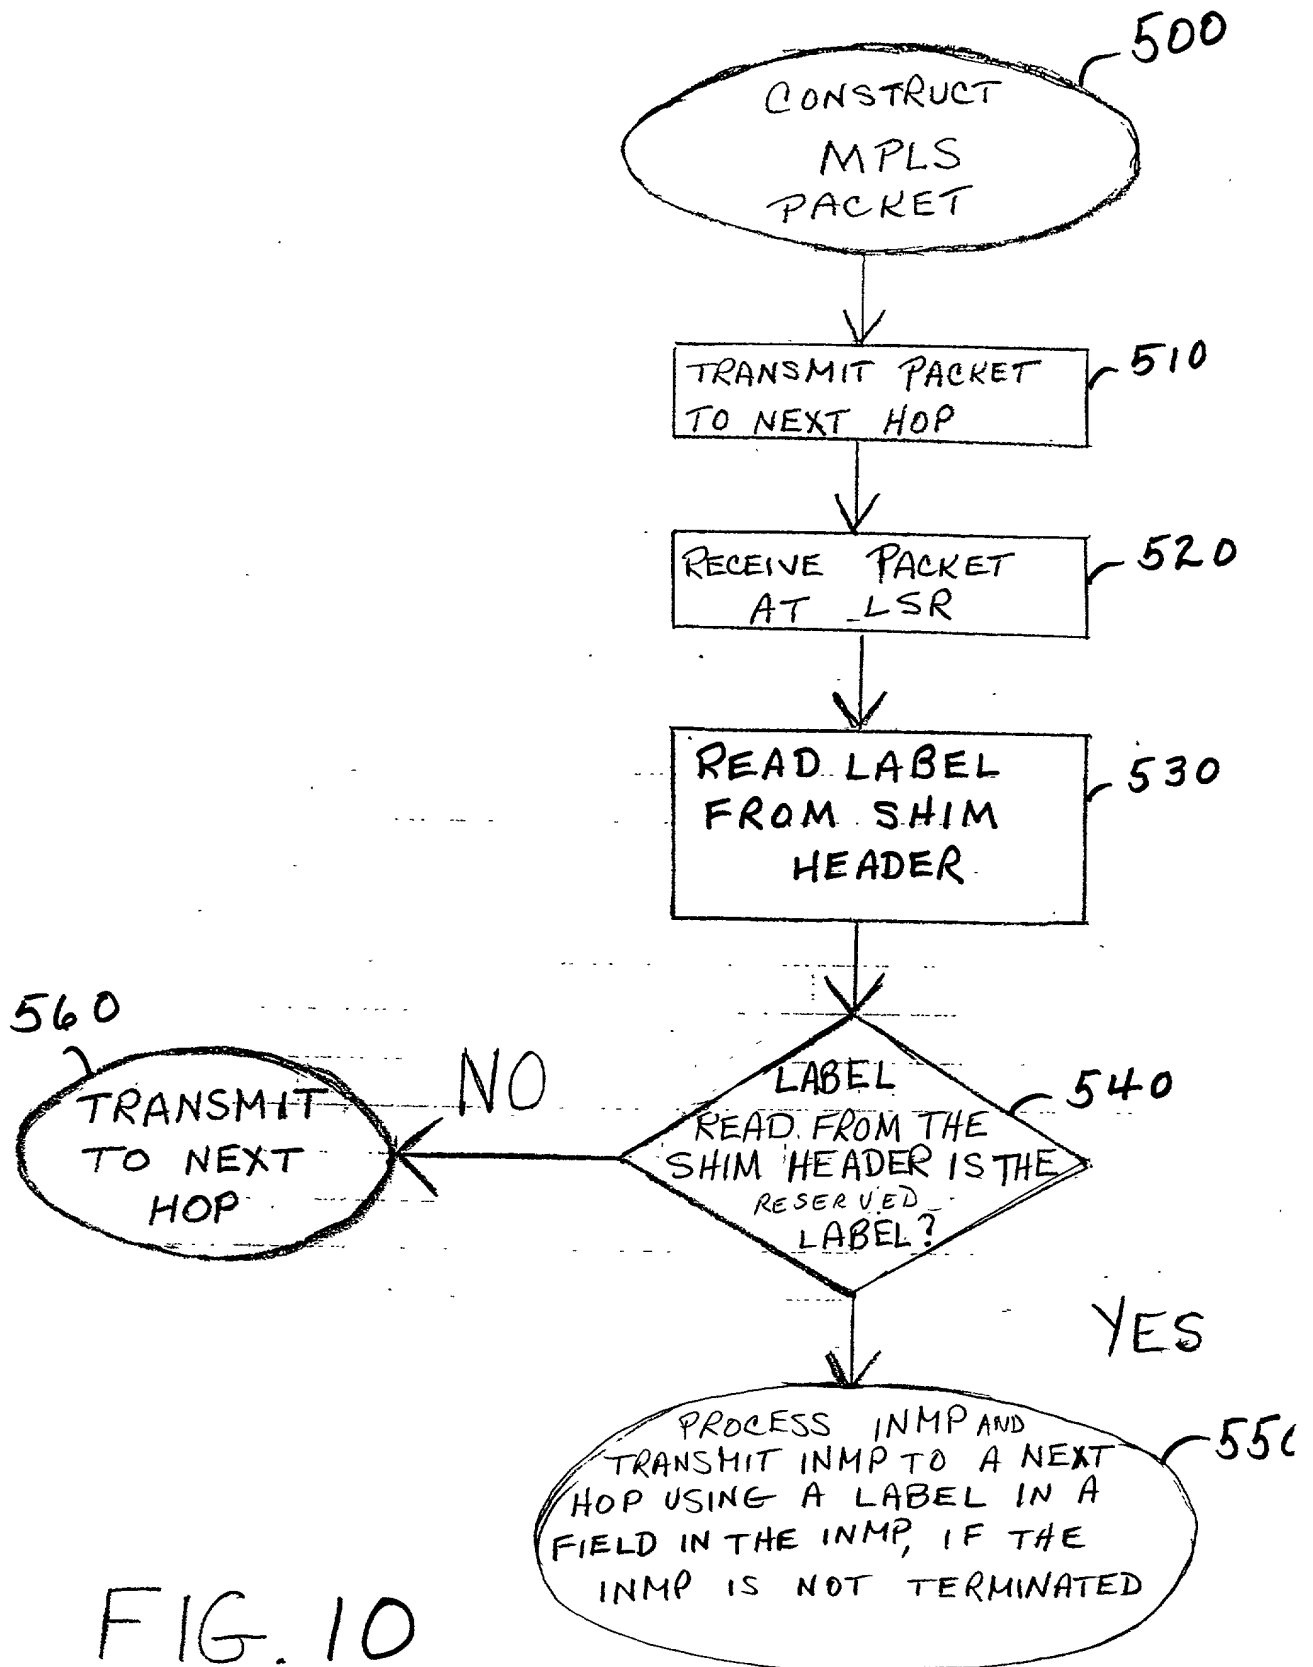


FIG. 10

002090 99460560

IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE

Declaration and Power of Attorney

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am an original, first and sole inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled **Techniques For Introducing In-Band Network Management Packets In Multi-Protocol Label Switching Networks**, the specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by an amendment, if any, specifically referred to in this oath or declaration.

I acknowledge the duty to disclose all information known to me which is material to patentability as defined in Title 37, Code of Federal Regulations, 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, 119 (a-d) or 365(a-b) of any PCT or foreign application(s) for patent or inventors' certificate listed below or priority benefits under 119(e) of any United States provisional application(s) listed below and have also identified below any foreign application for patent or inventors' certificate having a filing date before that of the application on which priority is claimed:

None

I hereby claim the benefit under Title 35, United States Code, 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, 112, we acknowledge the duty to disclose all information known to us to be material to patentability as defined in Title 37, Code of Federal Regulations, 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

None

002090" 93468560

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint the following attorney(s) with full power of substitution and revocation, to prosecute said application, to make alterations and amendments therein, to receive the patent, and to transact all business in the Patent and Trademark Office connected therewith:

Samuel H. Dworetsky	(Reg. No. 27873)
Thomas A. Restaino	(Reg. No. 33444)
Michele L. Conover	(Reg. No. 34962)
Cedric G. DeLaCruz	(Reg. No. 36498)
Rohini K. Garg	(Reg. No. 45272)
Benjamin S. Lee	(Reg. No. 42787)
Robert B. Levy	(Reg. No. 28234)
Susan E. McHale	(Reg. No. 35948)
Jeffrey M. Navon	(Reg. No. 32711)
Alfred G. Steinmetz	(Reg. No. 22971)

I also appoint Thomas H. Jackson (Reg. No. 29808), Frederic M. Meeker (Reg. No. 35282), and Ashok K. Mannava (Reg. No. 45301) of Banner & Witcoff as associate attorneys, with full power to prosecute said application, to make alterations and amendments therein, and to transact all business in the Patent and Trademark Office connected therewith.

Please address all correspondence to Mr. S. H. Dworetsky, AT&T Corp., P.O. Box 4110, Middletown, New Jersey 07748. Telephone calls should be made to Robert B. Levy at 908-221-5714.

Full name of sole inventor: Samson Boodaghians

Inventor's signature Samson Boodaghians Date 6/7/00

Residence: Wayside, Monmouth County, New Jersey

Citizenship: United States of America

Post Office Address: 33 Hillside Terrace
Wayside, New Jersey 07712

002090" 99768560